



Intel[®] Mobile Celeron[™] Processor Specification Update

Release Date: January 2000

Order Number: 244444-011

The Intel[®] Mobile Celeron[™] processor or the Intel[®] Celeron[™] Processor Mobile Module may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are documented in this Specification Update.

Information in this document is provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for particular purpose, merchantability or infringement or any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel® Mobile Celeron™ processor or the Intel® Celeron™ Processor Mobile Module may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

The Specification Update should be publicly available following the last shipment date for a period of time equal to the specific product's warranty period. Hardcopy Specification Updates will be available for one (1) year following End of Life (EOL). Web access will be available for three (3) years following EOL.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725 or by visiting Intel's website at <http://www.intel.com>
Copyright © Intel Corporation 1999, 2000.

*Third-party brands and names are the property of their respective owners

CONTENTS

REVISION HISTORY	ii
PREFACE	iii
Specification Update for the Intel® Mobile Celeron™ Processor	1
GENERAL INFORMATION	1
Intel® Mobile Celeron™ Processor (Micro-PGA) Markings	1
Intel® Mobile Celeron™ Processor (BGA) Markings	2
IDENTIFICATION INFORMATION	3
SUMMARY OF CHANGES	6
Summary of Errata	7
Summary of Documentation Changes	9
Summary of Specification Clarifications	10
Summary of Specification Changes	10
ERRATA	11
DOCUMENTATION CHANGES	39
SPECIFICATION CLARIFICATIONS	48
SPECIFICATION CHANGES	50

REVISION HISTORY

Date of Revision	Version	Description
January 1999	-001	This document is the first Specification Update for the Intel® Mobile Celeron™ processor.
March 1999	-002	Updated the Documentation Changes and Specifications Clarifications sections. Changed S-Spec Definition.
April 1999	-003	Added Erratum H43. Updated Processor Identification Information table and added footnote. Updated Processor Mobile Module Information table. Added MMC-2 marking diagram.
May 1999	-004	Updated Processor Identification Information table. Added Erratum H44. Added Documentation Changes H2 through H6. Added Specification Clarifications H1 and H2. Added Specification Change H1. Updated BGA1 marking diagram.
June 1999	-005	Updated the Intel Mobile Celeron Processor Identification Information table. Added μ PGA1 marking diagram. Added Erratum H45. Added Documentation Changes H7 through H10. Added Specification Change H2.
July 1999	-006	Added Erratum H46. Updated heading in Identification Information table. Updated the Intel Mobile Celeron Processor Identification and the Intel Celeron Processor Mobile Module Identification Information tables. Added Documentation Change H11. Added mcpA0 stepping to Summary Table of Changes. Updated Documentation Changes, Specification Clarifications, and Specification Changes introduction paragraphs.
August 1999	-007	Updated Preface, Documentation Changes, Specification Clarifications, and Specification Changes paragraphs. Added Documentation Change H12. Updated Codes Used in Summary Table. Updated column heading in Errata, Documentation Changes, Specification Clarifications and Specification Changes tables.
October 1999	-008	Added erratum H47
November 1999	-009	Added Errata H48, H49; added Documentation Change H13; added Specification Clarification H3; added Brand ID column in the <i>Identification Information</i> table. Updated the <i>Intel Mobile Celeron Processor Identification Information</i> table and notes. Updated the <i>Intel Celeron Processor Mobile Module Identification Information</i> table and notes. Added 433 MHz, and 466 MHz references.
December 1999	-010	Added Erratum H50 and Documentation Change H14
January 2000	-011	Added Erratum H51 and Documentation Change H15.

PREFACE

This document is an update to the specifications contained in the following documents:

- *Intel® Mobile Celeron™ Processor in Micro-PGA and BGA Package at 466 MHz, 433 MHz, 400 MHz, 366 MHz, 333 MHz, 300 MHz and 266 MHz* datasheet (Order Number 245112)
- *Intel® Mobile Celeron™ Processor: Mobile Module MMC-1* (Order Number 245101)
- *Intel® Mobile Celeron™ Processor: Mobile Module MMC-2* (Order Number 245102)
- *Intel Architecture Software Developer's Manual, Volumes 1, 2 and 3* (Order Numbers 243190, 243191, and 243192, respectively)

It is intended for hardware system manufacturers and software developers of applications, operating systems, or tools. It contains S-Specs, Errata, Documentation Changes, Specification Clarifications and, Specification Changes.

Nomenclature

S-Spec Number is a five-digit code used to identify products. Products are differentiated by their unique characteristics, e.g., core speed, L2 cache size, package type, etc. as described in the processor identification information table. Care should be taken to read all notes associated with each S-Spec number.

Errata are design defects or errors. Errata may cause the Intel Mobile Celeron processor or the Intel Mobile Celeron Module's behavior to deviate from published specifications. Hardware and software designed to be used with any given processor must assume that all errata documented for that processor are present on all devices unless otherwise noted.

Documentation Changes include typos, errors, or omissions from the current published specifications. These changes will be incorporated in the next release of the specifications.


Specification Clarifications describe a specification in greater detail or further highlight a specification's impact to a complex design situation. These clarifications will be incorporated in the next release of the specifications.

Specification Changes are modifications to the current published specifications for the Intel® Mobile Celeron™ processor or the Intel® Celeron™ Processor Mobile Module. These changes will be incorporated in the next release of the specifications.


Specification Update for the Intel® Mobile Celeron™ Processor

GENERAL INFORMATION

Intel® Mobile Celeron™ Processor (Micro-PGA) Markings



Intel® Mobile Celeron™ Processor in Micro-PGA Package

Intel® Mobile Celeron™ Processor (BGA) Markings**Intel® Mobile Celeron™ Processor in BGA1 Package**



IDENTIFICATION INFORMATION

The Intel Mobile Celeron processor or the Intel Celeron Processor Mobile Module can be identified by the following values:

Family ¹	266- , 300-, 333-, 366- , 400, 433, 466- Model 6 ²	Brand ID
0110	0110	00h (not supported)

NOTES:

1. The Family corresponds to bits [11:8] of the EDX register after RESET, bits [11:8] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the generation field of the Device ID register accessible through Boundary Scan.
2. The Model corresponds to bits [7:4] of the EDX register after RESET, bits [7:4] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the model field of the Device ID register accessible through Boundary Scan.

The Intel Mobile Celeron processor and the Intel Celeron Processor Mobile Module's second level (L2) cache size can be determined by the following register contents:

128-Kbyte Unified L2 Cache ¹	41h
---	-----

NOTES:

- 1 For the Intel® Mobile Celeron™ processor and the Intel® Celeron™ Processor Mobile Module, the unified L2 cache size will be returned as one of the cache/TLB descriptors when the CPUID instruction is executed with a 2 in the EAX register.

Intel® Mobile Celeron™ Processor Identification Information

S-Spec	Product Steppings	CPUID	Speed (MHz) Core/Bus	Integrated L2 Size (Kbytes)	Revision	Notes
SL23X	mcbA0	066Ah	233/66	128	BGA	1
SL23Y	mcbA0	066Ah	266/66	128	BGA	1
SL3AH	mcbA0	066Ah	300/66	128	BGA	1
SL3C8	mcbA0	066Ah	333/66	128	BGA	1
SL3C7	mcbA0	066Ah	366/66	128	BGA	1
SL3DQ	mcbA0	066Ah	266/66	128	BGA	2
SL3GQ	mcbA0	066Ah	400/66	128	BGA	1
SL3KA	mcbA0	066Ah	433/66	128	Micro-PGA	3
SL3KC	mcbA0	066Ah	466/66	128	Micro-PGA	3
SL3HM	mcpA0	066Ah	266/66	128	Micro-PGA	1
SL3HN	mcpA0	066Ah	300/66	128	Micro-PGA	1
SL3HP	mcpA0	066Ah	333/66	128	Micro-PGA	1
SL3HQ	mcpA0	066Ah	366/66	128	Micro-PGA	1
SL3GR	mcpA0	066Ah	400/66	128	Micro-PGA	1
SL3FL	mB0	128	0665h	466/66	PPGA	2
SL3KB	mcpA0	066Ah	433/66	128	Micro-PGA	3
SL3KD	mcpA0	066Ah	466/66	128	Micro-PGA	3

NOTES:

1. VCC_CORE is specified for 1.6 V +/-135 mV for these Intel® Mobile Celeron™ processors.
2. VCC_CORE is specified for 1.5 V +/-135 mV for these Intel Mobile Celeron processors.
3. VCC_CORE is specified for 1.9 V +/-135mV for these Intel Mobile Celeron processors.



Intel® Celeron™ Processor Mobile Module Identification Information

PTC	Product Steppings	CPUID	Speed (MHz) Core/Bus	Integrated L2 Size (Kbytes)	Package	Notes
PMH266010 01AA	cmmA0	066Ah	266/66	128	MMC-1	1
PMH300010 01AA	cmmA0	066Ah	300/66	128	MMC-1	1
PMH333010 01AA	cmmA0	066Ah	333/66	128	MMC-1	1
PMH366010 01AA	cmmA0	066Ah	366/66	128	MMC-1	1
PMH400010 01AA	cmmA0	066Ah	400/66	128	MMC-1	1
PMH433010 01AA	cmmA0	066Ah	433/66	128	MMC-1	2
PMH466010 01AA	cmmA0	066Ah	466/66	128	MMC-1	2
PMI2660100 1AA	cmmA0	066Ah	266/66	128	MMC-2	1
PMI3000100 1AA	cmmA0	066Ah	300/66	128	MMC-2	1
PMI3000100 1AA	cmmA0	066Ah	333/66	128	MMC-2	1
PMI3660100 1AA	cmmA0	066Ah	366/66	128	MMC-2	1
PMI4000100 1AA	cmmA0	066Ah	400/66	128	MMC-2	1
PMI4330100 1AA	cmmA0	066Ah	433/66	128	MMC-1	2
PMI4660100 1AA	cmmA0	066Ah	466/66	128	MMC-1	2

NOTES:

1. Vcore voltage is 1.6 V.
2. Vcore voltage is 1.9VV



SUMMARY OF CHANGES

The following table indicates the Errata, Documentation Changes, Specification Clarifications, or Specification Changes that apply to Intel® Mobile Celeron™ processors. Intel intends to fix some of the errata in a future stepping of the component, and to account for the other outstanding issues through documentation or specification changes as noted. This table uses the following notations:

CODES USED IN SUMMARY TABLE

X:	Erratum, Documentation Change, Specification Clarification or Specification Change applies to the given processor stepping.
(No mark) or (blank box):	This item is fixed in or does not apply to the given stepping.
Fix:	This erratum is intended to be fixed in a future stepping of the component.
Fixed:	This erratum has been previously fixed.
NoFix:	There are no plans to fix this erratum.
Doc:	Intel intends to update the appropriate documentation in a future revision.
PKG:	This column refers to errata on the Intel® Mobile Celeron™ processor or the Intel® Celeron™ Processor Mobile Module substrate.
AP:	APIC related erratum.
Shaded:	This item is either new or modified from the previous version of the document.

Each Specification Update item is prefixed with a capital letter to distinguish the product. The key below details the letters that are used in Intel's microprocessor Specification Updates:

A = Pentium® II processor

B = Mobile Pentium® II processor

C = Intel® Celeron™ processor

D = Pentium® II Xeon™ processor

E = Pentium® III processor

G = Pentium® III Xeon™ processor

H = Intel® Mobile Celeron™ processor

The Specification Updates for the Pentium® processor, Pentium® Pro processor, and other Intel products do not use this convention.

Summary of Errata

NO.	mcbA0	mcpA0	cmmA0	PKG	Plans	ERRATA
H1	X	X	X		NoFix	FP data operand pointer may be incorrectly calculated after FP access which wraps 64-Kbyte boundary in 16-bit code
H2	X	X	X		NoFix	Differences exist in debug exception reporting
H3	X	X	X		NoFix	Code fetch matching disabled debug register may cause debug exception
H4	X	X	X		NoFix	Double ECC error on read may result in BINIT#
H5	X	X	X		NoFix	FP inexact-result exception flag may not be set
H6	X	X	X		NoFix	BTM for SMI will contain incorrect FROM EIP
H7	X	X	X		NoFix	I/O restart in SMM may fail after simultaneous MCE
H8	X	X	X		NoFix	Branch traps do not function if BTMs are also enabled
H9	X	X	X		NoFix	Machine check exception handler may not always execute successfully
H10	X	X	X		NoFix	MCE due to L2 parity error gives L1 MCACOD.LL
H11	X	X	X		NoFix	LBERR may be corrupted after some events
H12	X	X	X		NoFix	BTMs may be corrupted during simultaneous L1 cache line replacement
H13	X	X	X		Fix	Potential early deassertion of LOCK# during split-lock cycles
H14	X	X	X		NoFix	A20M# may be inverted after returning from SMM and Reset
H15	X	X	X		Fix	Reporting of floating-point exception may be delayed
H16	X	X	X		NoFix	Near CALL to ESP creates unexpected EIP address
H17	X	X	X		Fix	Built-in self test always gives nonzero result
H18	X	X	X		Fix	Cache state corruption in the presence of page A/D-bit setting and snoop traffic
H19	X	X	X		Fix	Snoop cycle generates spurious machine check exception
H20	X	X	X		Fix	MOVD/MOVB instruction writes to memory prematurely



Summary of Errata

NO.	mcbA0	mcpA0	cmmA0	PKG	Plans	ERRATA
H21	X	X	X		NoFix	Memory type undefined for nonmemory operations
H22	X	X	X		NoFix	FP data operand pointer may not be zero after power on or reset
H23	X	X	X		NoFix	MOVD following zeroing instruction can cause incorrect result
H24	X	X	X		NoFix	Premature execution of a load operation prior to exception handler invocation
H25	X	X	X		NoFix	Read portion of RMW instruction may execute twice
H26	X	X	X		Fix	Intervening writeback may occur during locked transaction
H27	X	X	X		NoFix	MC2_STATUS MSR has model-specific error code and machine check architecture error code reversed
H28	X	X	X		NoFix	Mixed cacheability of lock variables is problematic in MP systems
H29			X		Fix	Thermal sensor may assert SMBALERT# incorrectly
H30	X	X	X		NoFix	MOV with debug register causes debug exception
H31	X	X	X		NoFix	Upper four PAT entries not usable with Mode B or Mode C paging
H32	X	X	X		Fix	Incorrect memory type may be used when MTRRs are disabled
H33	X	X	X		Fix	Misprediction in program flow may cause unexpected instruction execution
H34	X	X	X		NoFix	Data breakpoint exception in a displacement relative near call may corrupt EIP
H35	X	X	X		NoFix	System bus ECC not functional with 2:1 ratio
H36	X	X	X		Fix	Fault on REP CMPS/SCAS operation may cause incorrect EIP
H37	X	X	X		NoFix	RDMSR and WRMSR to invalid MSR address may not cause GP fault
H38	X	X	X		NoFix	SYSENTER/SYSEXIT instructions can implicitly load "null segment selector" to SS and CS registers
H39	X	X	X		NoFix	PRELOAD followed by EXTEST does not load boundary scan data
H40	X	X	X		NoFix	Far jump to new TSS with D-bit cleared may

Summary of Errata

NO.	mcbA0	mcpA0	cmmA0	PKG	Plans	ERRATA
						cause system hang
H41	X	X	X		NoFix	Incorrect chunk ordering may prevent execution of the Machine Check Exception handler after BINIT#
H42	X	X	X		NoFix	Resume Flag may not be cleared after debug exception
H43			X		Fix	Processor may return invalid parameters on execution of the CPUID instruction
H44	X	X	X		NoFix	Internal cache protocol violation may cause system hang
H45	X	X	X		NoFix	GP# fault on WSMSR to ROB_CR_BKUPTMPDR6
H46	X	X	X		NoFix	Machine check exception may occur due to improper line eviction in the IFU
H47	X	X	X		NoFix	Lower Bits of SMRAM SMBASE Register Cannot Be Written With an ITP
H48	X	X	X		NoFix	Task switch caused by page fault may cause wrong PTE and PDE access bit to be set
H49	X	X	X		NoFix	Cross modifying code operations on a jump instruction may cause a general protection fault
H50	X	X	X		NoFix	Deadlock may occur due to illegal-instruction/page-miss combination
H51	X	X	X		NoFix	FLUSH# assertion following STPCLK# may prevent CPU clocks from stopping

Summary of Documentation Changes

NO.	mcbA0	mcpA0	cmmA0	PKG	Plans	DOCUMENTATION CHANGES
H1	X	X	X		Doc	SMBus data setup time
H2			X		Doc	Intel® Celeron™ Processor Mobile Module Connector 1 (MMC-1) northbridge correction
H3			X		Doc	VR_ON and VRPWRGD specification correction for MMC-1 and MMC-2
H4			X		Doc	STR and POS measurement correction for MMC-1 and MMC-2
H5			X		Doc	L2 cache correction for MMC-1 and MMC-2
H6	X	X			Doc	Die height (A ₂) in surface-mount BGA1

Summary of Documentation Changes

NO.	mcbA0	mcpA0	cmmA0	PKG	Plans	DOCUMENTATION CHANGES
						package specification
H7	X	X	X		Doc	MTRR initialization clarification
H8	X	X	X		Doc	STPCLK# pin definition
H9	X	X	X		Doc	RESET# pin definition
H10	X				Doc	BGA1 package mechanical specifications
H11	X	X			Doc	Generic clock waveform figure
H12	X	X	X		Doc	Invalidating caches and TLBs
H13	X	X	X		Doc	Handling of self-modifying and cross-modifying code
H14	X	X	X		Doc	Machine check architecture initialization of MCi_STATUS registers
H15	X	X	X		Doc	LOCK# signal prefix operands

Summary of Specification Clarifications

NO.	mcbA0	mcpA0	cmmA0	PKG	Plans	SPECIFICATION CLARIFICATIONS
H1	X	X			Doc	Shipping container maximum temperature rating for BGA1
H2			X		Doc	ESD for MMC-1 and MMC-2
H3			X		Doc	Bulk capacitance requirements for the system electronics

Summary of Specification Changes

NO.	mcbA0	mcpA0	cmmA0	PKG	Plans	SPECIFICATION CHANGES
H1	X	X			Doc	I _{CC,SG} , I _{CC,QS} , and I _{CC,DSL} P maximum specifications for BGA1 and mini-cartridge
H2	X	X			Doc	Temperature note for thermal power specifications

ERRATA

H1. FP Data Operand Pointer May Be Incorrectly Calculated After FP Access Which Wraps 64-Kbyte Boundary in 16-Bit Code

Problem: The FP Data Operand Pointer is the effective address of the operand associated with the last noncontrol floating-point instruction executed by the machine. If an 80-bit floating-point access (load or store) occurs in a 16-bit mode other than protected mode (in which case the access will produce a segment limit violation), the memory access wraps a 64-Kbyte boundary, and the floating-point environment is subsequently saved, the value contained in the FP Data Operand Pointer may be incorrect.

Implication: A 32-bit operating system running 16-bit floating-point code may encounter this erratum, under the following conditions:

- The operating system is using a segment greater than 64 Kbytes in size.
- An application is running in a 16-bit mode other than protected mode.
- An 80-bit floating-point load or store which wraps the 64-Kbyte boundary is executed.
- The operating system performs a floating-point environment store (FSAVE/FNSAVE/STENV/FNSTENV) after the above memory access.
- The operating system uses the value contained in the FP Data Operand Pointer.

Wrapping an 80-bit floating-point load around a segment boundary in this way is not a normal programming practice. Intel has not currently identified any software which exhibits this behavior.

Workaround: If the FP Data Operand Pointer is used in an OS which may run 16-bit floating-point code, care must be taken to ensure that no 80-bit floating-point accesses are wrapped around a 64-Kbyte boundary.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H2. Differences Exist in Debug Exception Reporting

Problem: There exist some differences in the reporting of code and data breakpoint matches between that specified by previous Intel processors' specifications and the behavior of the Intel Mobile Celeron processor, as described below:

Case 1: The first case is for a breakpoint set on a MOVSS or POPSS instruction, when the instruction following it causes a debug register protection fault (DR7.gd is already set, enabling the fault). The processor reports delayed data breakpoint matches from the MOVSS or POPSS instructions by setting the matching DR6.bi bits, along with the debug register protection fault (DR6.bd). If additional breakpoint faults are matched during the call of the debug fault handler, the processor sets the breakpoint match bits (DR6.bi) to reflect the breakpoints matched by both the MOVSS or POPSS breakpoint and the debug fault handler call. The Intel Mobile Celeron processor only sets DR6.bd in either situation, and does not set any of the DR6.bi bits.

Case 2: In the second breakpoint reporting failure case, if a MOVSS or POPSS instruction with a data breakpoint is followed by a store to memory which crosses a 4- Kbyte page boundary, the breakpoint information for the MOVSS or POPSS will be lost. Previous processors retain this information across such a page split.

Case 3: If they occur after a MOVSS or POPSS instruction, the INT *n*, INTO, and INT3 instructions zero the DR6.bi bits (bits B0 through B3), clearing pending breakpoint information, unlike previous processors.

Case 4: If a data breakpoint and an SMI (System Management Interrupt) occur simultaneously, the SMI will be serviced via a call to the SMM handler, and the pending breakpoint will be lost.

Case 5: When an instruction which accesses a debug register is executed, and a breakpoint is encountered on the instruction, the breakpoint is reported twice.

Implication: When debugging or when developing debuggers for a Intel Mobile Celeron processor-based system, this behavior should be noted. Normal usage of the MOVSS or POPSS instructions (e.g., following them with a MOV ESP) will not exhibit the behavior of cases 1-3. Debugging in conjunction with SMM will be limited by case 4.

Workaround: Following MOVSS and POPSS instructions with a MOV ESP instruction when using breakpoints will avoid the first three cases of this erratum. No workaround has been identified for cases 4 or 5.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H3. Code Fetch Matching Disabled Debug Register May Cause Debug Exception

Problem: The bits L0-3 and G0-3 enable breakpoints local to a task and global to all tasks, respectively. If one of these bits is set, a breakpoint is enabled, corresponding to the addresses in the debug registers DR0 - DR3. If at least one of these breakpoints is enabled, any of these registers are *disabled* (e.g., Ln and Gn are 0), and RWN for the disabled register is 00 (indicating a breakpoint on instruction execution), normally an instruction fetch will not cause an instruction-breakpoint fault based on a match with the address in the disabled register(s). However, if the address in a disabled register matches the address of a code fetch which also results in a page fault, an instruction-breakpoint fault will occur.

Implication: While debugging software, extraneous instruction-breakpoint faults may be encountered if breakpoint registers are not cleared when they are disabled. Debug software which does not implement a code breakpoint handler will fail, if this occurs. If a handler is present, the fault will be serviced. Mixing data and code may exacerbate this problem by allowing disabled data breakpoint registers to break on an instruction fetch.

Workaround: The debug handler should clear breakpoint registers before they become disabled.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H4. Double ECC Error on Read May Result in BINIT#

Problem: For this erratum to occur, the following conditions must be met:

- Machine Check Exceptions (MCEs) must be enabled.
- A dataless transaction (such as a write invalidate) must be occurring simultaneously with a transaction which returns data (a normal read).
- The read data must contain a double-bit uncorrectable ECC error.

If these conditions are met, the Intel Mobile Celeron processor will not be able to determine which transaction was erroneous, and instead of generating an MCE, it will generate a BINIT#.

Implication: The bus will be reinitialized in this case. However, since a double-bit uncorrectable ECC error occurred on the read, the MCE handler (which is normally reached on a double-bit uncorrectable ECC error for a read) would most likely cause the same BINIT# event.

Workaround: Though the ability to drive BINIT# can be disabled in the Intel Mobile Celeron processor, which would prevent the effects of this erratum, overall system behavior would not improve, since the error which would normally cause a BINIT# would instead cause the machine to shut down. No other workaround has been identified.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.



H5. FP Inexact-Result Exception Flag May Not Be Set

Problem: When the result of a floating-point operation is not exactly representable in the destination format (1/3 in binary form, for example), an inexact-result (precision) exception occurs. When this occurs, the PE bit (bit 5 of the FPU status word) is normally set by the processor. Under certain rare conditions, this bit may not be set when this rounding occurs. However, other actions taken by the processor (invoking the software exception handler if the exception is unmasked) are not affected. This erratum can only occur if the floating-point operation which causes the precision exception is immediately followed by one of the following instructions:

- FST m32real
- FST m64real
- FSTP m32real
- FSTP m64real
- FSTP m80real
- FIST m16int
- FIST m32int
- FISTP m16int
- FISTP m32int
- FISTP m64int

Note that even if this combination of instructions is encountered, there is also a dependency on the internal pipelining and execution state of both instructions in the processor.

Implication: Inexact-result exceptions are commonly masked or ignored by applications, as it happens frequently, and produces a rounded result acceptable to most applications. The PE bit of the FPU status word may not always be set upon receiving an inexact-result exception. Thus, if these exceptions are unmasked, a floating-point error exception handler may not recognize that a precision exception occurred. Note that this is a “sticky” bit, e.g., once set by an inexact-result condition, it remains set until cleared by software.

Workaround: This condition can be avoided by inserting two NOP instructions between the two floating-point instructions.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H6. BTM for SMI Will Contain Incorrect FROM EIP

Problem: A system management interrupt (SMI) will produce a Branch Trace Message (BTM), if BTMs are enabled. However, the FROM EIP field of the BTM (used to determine the address of the instruction which was being executed when the SMI was serviced) will not have been updated for the SMI, so the field will report the same FROM EIP as the previous BTM.

Implication: A BTM which is issued for an SMI will not contain the correct FROM EIP, limiting the usefulness of BTMs for debugging software in conjunction with System Management Mode (SMM).

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H7. I/O Restart in SMM May Fail After Simultaneous MCE

Problem: If an I/O instruction (IN, INS, REP INS, OUT, OUTS, or REP OUTS) is being executed, and if the data for this instruction becomes corrupted, the Intel Mobile Celeron processor will signal a machine check exception (MCE). If the instruction is directed at a device which is powered down, the processor may also receive an assertion of SMI#. Since MCEs have higher priority, the processor will call the MCE handler, and the SMI# assertion will remain pending. However, upon attempting to execute the first instruction of the MCE handler, the SMI# will be recognized and the processor will attempt to execute the SMM handler. If the SMM handler is completed successfully, it will attempt to restart the I/O instruction, but will not have the correct machine state, due to the call to the MCE handler.

Implication: A simultaneous MCE and SMI# assertion may occur for one of the I/O instructions above. The SMM handler may attempt to restart such an I/O instruction, but will have corrupted state due to the MCE handler call, leading to failure of the restart and shutdown of the processor.

Workaround: If a system implementation must support both SMM and MCEs, the first thing the SMM handler code (when an I/O restart is to be performed) should do is check for a pending MCE. If there is an MCE pending, the SMM handler should immediately exit via an RSM instruction and allow the machine check exception handler to execute. If there is not, the SMM handler may proceed with its normal operation.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H8. Branch Traps Do Not Function if BTMs Are Also Enabled

Problem: If branch traps or branch trace messages (BTMs) are enabled alone, both function as expected. However, if both are enabled, only the BTMs will function, and the branch traps will be ignored.

Implication: The branch traps and branch trace message debugging features cannot be used together.

Workaround: If branch trap functionality is desired, BTMs must be disabled.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H9. Machine Check Exception Handler May Not Always Execute Successfully

Problem: An asynchronous machine check exception (MCE), such as a BINIT# event, which occurs during an access that splits a 4-Kbyte page boundary may leave some internal registers in an indeterminate state. Thus, MCE handler code may not always run successfully if an asynchronous MCE has occurred previously.

Implication: An MCE may not always result in the successful execution of the MCE handler. However, asynchronous MCEs usually occur upon detection of a catastrophic system condition that would also hang the processor. Leaving MCEs disabled will result in the condition which caused the asynchronous MCE instead causing the processor to enter shutdown. Therefore, leaving MCEs disabled may not improve overall system behavior.

Workaround: No workaround which would guarantee successful MCE handler execution under this condition has been identified.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H10. MCE Due to L2 Parity Error Gives L1 MCACOD.LL

Problem: If a Cache Reply Parity (CRP) error, Cache Address Parity (CAP) error, or Cache Synchronous Error (CSER) occurs on an access to the Intel Mobile Celeron processor's L2 cache, the resulting Machine Check Architectural Error Code (MCACOD) will be logged with '01' in the LL field. This value indicates an L1 cache error; the value should be '10', indicating an L2 cache error. Note that L2 ECC errors have the correct value of '10' logged.

Implication: An L2 cache access error, other than an ECC error, will be improperly logged as an L1 cache error in MCACOD.LL.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H11. LBER May Be Corrupted After Some Events

Problem: The last branch record (LBR) and the last branch before exception record (LBER) can be used to determine the source and destination information for previous branches or exceptions. The LBR contains the source and destination addresses for the last branch or exception, and the LBER contains similar information for the last branch taken before the last exception. This information is typically used to determine the location of a branch which leads to execution of code which causes an exception. However, after a catastrophic bus condition which results in an assertion of BINIT# and the reinitialization of the buses, the value in the LBER may be corrupted. Also, after either a CALL which results in a fault or a software interrupt, the LBER and LBR will be updated to the same value, when the LBER should not have been updated.

Implication: The LBER and LBR registers are used only for debugging purposes. When this erratum occurs, the LBER will not contain reliable address information. The value of LBER should be used with caution when debugging branching code; if the values in the LBR and LBER are the same, then the LBER value is incorrect. Also, the value in the LBER should not be relied upon after a BINIT# event.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H12. BTMs May Be Corrupted During Simultaneous L1 Cache Line Replacement

Problem: When Branch Trace Messages (BTMs) are enabled and such a message is generated, the BTM may be corrupted when issued to the bus by the L1 cache if a new line of data is brought into the L1 data cache simultaneously. Though the new line being stored in the L1 cache is stored correctly, and no corruption occurs in the data, the information in the BTM may be incorrect due to the internal collision of the data line and the BTM.

Implication: Although BTMs may not be entirely reliable due to this erratum, the conditions necessary for this boundary condition to occur have only been exhibited during focused simulation testing. Intel has currently not observed this erratum in a system level validation environment.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H13. Potential Early Deassertion of LOCK# During Split-Lock Cycles

Problem: During a split-lock cycle there are four bus transactions: 1st ADS# (a partial read), 2nd ADS# (a partial read), 3rd ADS# (a partial write), and the 4th ADS# (a partial write). Due to this erratum, LOCK# may deassert one clock after the 4th ADS# of the split-lock cycle instead of after the 4th RS# assertion corresponding to the 4th ADS# has been sampled. The following sequence of events are required for this erratum to occur:

1. A lock cycle occurs (split or nonsplit).
2. Five more bus transactions (assertion of ADS#) occur.
3. A split-lock cycle occurs and BNR# toggles after the 3rd ADS# (partial write) of the split-lock cycle. This in turn delays the assertion of the 4th ADS# of the split-lock cycle. BNR# toggling at this time could most likely happen when the bus is set for an IOQ depth of 2.

When all of these events occur, LOCK# will be deasserted in the next clock after the 4th ADS# of the split-lock cycle.

Implication: This may affect chipset logic which monitors the behavior of LOCK# deassertion.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H14. A20M# May Be Inverted After Returning From SMM and Reset

Problem: This erratum is seen when software causes the following events to occur:

1. The assertion of A20M# in real address mode.
2. After entering the 1-Mbyte address wrap-around mode caused by the assertion of A20M#, there is an assertion of SMI# intended to cause a Reset or remove power to the processor. Once in the SMM handler, software saves the SMM state save map to an area of nonvolatile memory from which it can be restored at some point in the future. Then software asserts RESET# or removes power to the processor.
3. After exiting Reset or completion of power-on, software asserts SMI# again. Once in the SMM handler, it then retrieves the old SMM state save map which was saved in event 2 above and copies it into the current SMM state save map. Software then asserts A20M# and executes the RSM instruction. After exiting the SMM handler, the polarity of A20M# is inverted.

Implication: If this erratum occurs, A20M# will behave with a polarity opposite from what is expected (e.g., the 1-Mbyte address wrap-around mode is enabled when A20M# is deasserted, and does not occur when A20M# is asserted).

Workaround: Software should save the A20M# signal state in nonvolatile memory before an assertion of RESET# or a power down condition. After coming out of Reset or at power on, SMI# should be asserted again. During the restoration of the old SMM state save map described in event 3 above, the entire map should be restored, except for bit 5 of the byte at offset 7F18h. This bit should retain the value assigned to it when the SMM state save map was created in event 3. The SMM handler should then restore the original value of the A20M# signal.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H15. Reporting of Floating-Point Exception May Be Delayed

Problem: The Intel Mobile Celeron processor normally reports a floating-point exception for an instruction when the next floating-point or MMX™ technology instruction is executed. The assertion of FERR# and/or the INT 16 interrupt corresponding to the exception may be delayed until the floating-point or MMX technology instruction *after* the one which is expected to trigger the exception, if the following conditions are met:

1. A floating-point instruction causes an exception.
2. Before another floating-point or MMX technology instruction, any one of the following occurs:
 - a. A subsequent data access occurs to a page which has not been marked as accessed
 - b. Data is referenced which crosses a page boundary, or
 - c. A possible page-fault condition is detected which, when resolved, completes without faulting.
3. The instruction causing event 2 above is followed by a MOVQ or MOVD store instruction.

Implication: This erratum only affects software which operates with floating-point exceptions unmasked. Software which requires floating-point exceptions to be visible on the next floating-point or MMX technology instruction, and which uses floating-point calculations on data which is then used for MMX technology instructions, may see a delay in the reporting of a floating-point instruction exception in some cases. Note that mixing floating-point and MMX technology instructions in this way is not recommended.

Workaround: Inserting a WAIT or FWAIT instruction (or reading the floating-point status register) between the floating-point instruction and the MOVQ or MOVD instruction will give the expected results. This is already the recommended practice for software.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H16. Near CALL to ESP Creates Unexpected EIP Address

Problem: As documented, the CALL instruction saves procedure linking information in the procedure stack and jumps to the called procedure specified with the destination (target) operand. The target operand specifies the address of the first instruction in the called procedure. This operand can be an immediate value, a general purpose register, or a memory location. When accessing an absolute address indirectly using the stack pointer (ESP) as a base register, the base value used is the value in the ESP register before the instruction executes. However, when accessing an absolute address directly using ESP as the base register, the base value used is the value of ESP *after* the return value is pushed on the stack, not the value in the ESP register *before* the instruction executed.

Implication: Due to this erratum, the processor may transfer control to an unintended address. Results are unpredictable, depending on the particular application, and can range from no effect to the unexpected termination of the application due to an exception. Intel has observed this erratum only in a focused testing environment. Intel has not observed any commercially available operating system, application, or compiler that makes use of or generates this instruction.

Workaround: If the other seven general purpose registers are unavailable for use, and it is necessary to do a CALL via the ESP register, first push ESP onto the stack, then perform an *indirect* call using ESP (e.g., CALL [ESP]). The saved version of ESP should be popped off the stack after the call returns.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H17. Built-in Self Test Always Gives Nonzero Result

Problem: The Built-in Self Test (BIST) of the Intel Mobile Celeron processor does not give a zero result to indicate a passing test. Regardless of pass or fail status, bit 6 of the BIST result in the EAX register after running BIST is set.

Implication: Software which relies on a zero result to indicate a passing BIST will indicate BIST failure.

Workaround: Mask bit 6 of the BIST result register when analyzing BIST results.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H18. Cache State Corruption in the Presence of Page A/D-bit Setting and Snoop Traffic

Problem: If an operating system uses the Page Access and/or Dirty bit feature implemented in the Intel architecture and there is a significant amount of snoop traffic on the bus, while the processor is setting the Access and/or Dirty bit the processor may inappropriately change a single L1 cache line to the modified state.

Implication: The occurrence of this erratum may result in cache incoherency, which may cause parity errors, data corruption (with no parity error), unexpected application or operating system termination, or system hangs.

Workaround: It is possible for BIOS code to contain a workaround for this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.



H19. Snoop Cycle Generates Spurious Machine Check Exception

Problem: The processor may incorrectly generate a Machine Check Exception (MCE) when it processes a snoop access that does not hit the L1 data cache. Due to an internal logic error, this type of snoop cycle may still check data parity on undriven data lines. The processor generates a spurious machine check exception as a result of this unnecessary parity check.

Implication: A spurious machine check exception may result in an unexpected system halt if Machine Check Exception reporting is enabled in the operating system.

Workaround: It is possible for BIOS code to contain a workaround for this erratum. This workaround would fix the erratum; however, the data parity error will still be reported.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H20. MOVD/MOVQ Instruction Writes to Memory Prematurely

Problem: When an instruction encounters a fault, the faulting instruction should not modify any CPU or system state. However, when the MMX™ technology store instructions MOVD and MOVQ encounter any of the following events, it is possible for the store to be committed to memory even though it should be canceled:

1. If CR0.EM = 1 (Emulation bit), then the store could happen prior to the triggered invalid opcode exception.
2. If the floating-point Top-of-Stack (FP TOS) is not zero, then the store could happen prior to executing the processor assist routine that sets the FP TOS to zero.
3. If there is an unmasked floating-point exception pending, then the store could happen prior to the triggered unmasked floating-point exception.
4. If CR0.TS = 1 (Task Switched bit), then the store could happen prior to the triggered Device Not Available (DNA) exception.

If the MOVD/MOVQ instruction is restarted after handling any of the above events, then the store will be performed again, overwriting with the expected data. The instruction will not be restarted after event 1. The instruction will definitely be restarted after events 2 and 4. The instruction may or may not be restarted after event 3, depending on the specific exception handler.

Implication: This erratum causes unpredictable behavior in an application if MOVD/MOVQ instructions are used to manipulate semaphores for multiprocessor synchronization, or if these MMX instructions are used to write to uncacheable memory or memory mapped I/O that has side effects, e.g., graphics devices. This erratum is completely transparent to all applications that do not have these characteristics. When each of the above conditions are analyzed:

1. Setting the CR0.EM bit forces all floating-point/MMX instructions to be handled by software emulation. The MOVD/MOVQ instruction, which is an MMX instruction, would be considered an invalid instruction. Operating systems typically terminates the application after getting the expected invalid opcode fault.
2. The FP TOS not equal to 0 case only occurs when the MOVD/MOVQ store is the first MMX instruction in an MMX technology routine and the previous floating-point routine did not clean up the floating-point states properly when it exited. Floating-point routines commonly leave TOS to 0 prior to exiting. For a store to be executed as the first MMX instruction in an MMX technology routine following a floating-point routine, the software would be implementing instruction level intermixing of floating-point and MMX instructions. Intel does not recommend this practice.

3. The unmasked floating-point exception case only occurs if the store is the first MMX technology instruction in an MMX technology routine and the previous floating-point routine exited with an unmasked floating-point exception pending. Again, for a store to be executed as the first MMX instruction in an MMX technology routine following a floating-point routine, the software would be implementing instruction level intermixing of floating-point and MMX instructions. Intel does not recommend this practice.
4. Device Not Available (DNA) exceptions occur naturally when a task switch is made between two tasks that use either floating-point instructions and/or MMX instructions. For this erratum, in the event of the DNA exception, data from the prior task may be temporarily stored to the present task's program state.

Workaround: Do not use MMX instructions to manipulate semaphores for multiprocessor synchronization. Do not use MOVD/MOVQ instructions to write directly to I/O devices if doing so triggers user visible side effects. An OS can prevent old data from being stored to a new task's program state by cleansing the FPU explicitly after every task switch. Follow Intel's recommended programming paradigms in the *Intel Architecture Optimization Manual* for writing MMX technology programs. Specifically, do not mix floating-point and MMX instructions. When transitioning to new a MMX technology routine, begin with an instruction that does not depend on the prior state of either the MMX technology registers or the floating-point registers, such as a load or PXOR mm0, mm0. Be sure that the FP TOS is clear before using MMX instructions.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H21. Memory Type Undefined for Nonmemory Operations

Problem: The Memory Type field for nonmemory transactions such as I/O and Special Cycles are undefined. Although the Memory Type attribute for nonmemory operations logically should (and usually does) manifest itself as UC, this feature is not designed into the implementation and is therefore inconsistent.

Implication: Bus agents may decode a non-UC memory type for nonmemory bus transactions.

Workaround: Bus agents must consider transaction type to determine the validity of the Memory Type field for a transaction.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H22. FP Data Operand Pointer May Not Be Zero After Power On or Reset

Problem: The FP Data Operand Pointer, as specified, should be reset to zero upon power on or Reset by the processor. Due to this erratum, the FP Data Operand Pointer may be nonzero after power on or Reset.

Implication: Software which uses the FP Data Operand Pointer and count on its value being zero after power on or Reset without first executing an FINIT/FNINIT instruction will use an incorrect value, resulting on incorrect behavior of the software.

Workaround: Software should follow the recommendation in Section 8.2 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide* (Order Number 243192). This recommendation states that if the FPU will be used, software-initialization code should execute an FINIT/FNINIT instruction following a hardware reset. This will correctly clear the FP Data Operand Pointer to zero.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H23. MOVD Following Zeroing Instruction Can Cause Incorrect Result

Problem: An incorrect result may be calculated after the following circumstances occur:

1. A register has been zeroed with either a SUB reg, reg instruction or an XOR reg, reg instruction
2. A value is moved with sign extension into the same register's lower 16 bits; or a signed integer multiply is performed to the same register's lower 16 bits
3. This register is then copied to an MMX™ technology register using the MOVD instruction prior to any other operations on the sign-extended value.

Specifically, the sign may be incorrectly extended into bits 16-31 of the MMX technology register. Only the MMX technology register is affected by this erratum.

The erratum only occurs when the 3 following steps occur in the order shown. The erratum may occur with up to 40 intervening instructions that do not modify the sign-extended value between steps 2 and 3.

1. XOR EAX, EAX
or SUB EAX, EAX
2. MOVSBX AX, BL
or MOVSBX AX, byte ptr <memory address> or MOVSBX AX, BX
or MOVSBX AX, word ptr <memory address> or IMUL BL (AX implicit, opcode F6 /5)
or IMUL byte ptr <memory address> (AX implicit, opcode F6 /5) or IMUL AX, BX (opcode 0F AF /r)
or IMUL AX, word ptr <memory address> (opcode 0F AF /r) or IMUL AX, BX, 16 (opcode 6B /r ib)
or IMUL AX, word ptr <memory address>, 16 (opcode 6B /r ib) or IMUL AX, 8 (opcode 6B /r ib)
or IMUL AX, BX, 1024 (opcode 69 /r iw)
or IMUL AX, word ptr <memory address>, 1024 (opcode 69 /r iw) or IMUL AX, 1024 (opcode 69 /r iw)
or CBW
3. MOVD MM0, EAX

Note that the values for immediate byte/words are merely representative (i.e., 8, 16, 1024) and that any value in the range for the size may be affected. Also, note that this erratum may occur with "EAX" replaced with any 32-bit general purpose register, and "AX" with the corresponding 16-bit version of that replacement. "BL" or "BX" can be replaced with any 8-bit or 16-bit general purpose register. The CBW and IMUL (opcode F6 /5) instructions are specific to the EAX register only.

In the example, EAX is forced to contain 0 by the XOR or SUB instructions. Since the MOVSBX, IMUL and CBW instructions listed should modify only bits 15:8 of EAX by sign extension, bits 31:16 of EAX should always contain 0. This implies that when MOVD copies EAX to MM0, bits 31:16 of MM0 should also be 0. Under certain scenarios, bits 31:16 of MM0 are not 0, but are replicas of bit 15 (the 16th bit) of AX. This is noticeable when the value in AX after the MOVSBX, IMUL or CBW instruction is negative, i.e., bit 15 of AX is a 1.

When AX is positive (bit 15 of AX is a 0), MOVD will always produce the correct answer. If AX is negative (bit 15 of AX is a 1), MOVD may produce the right answer or the wrong answer depending on the point in time when the MOVD instruction is executed in relation to the MOVSBX, IMUL or CBW instruction.

Implication: The effect of incorrect execution will vary from unnoticeable, due to the code sequence discarding the incorrect bits, to an application failure. If the MMX technology-enabled application in which MOVD is used to manipulate pixels, it is possible for one or more pixels to exhibit the wrong color or position momentarily. It is also possible for a computational application that uses the MOVD instruction in the manner described above to produce incorrect data. Note that this data may cause an unexpected page fault or general protection fault.

Workaround: There are two possible workarounds for this erratum:

1. Rather than using the MOVXS-MOVD or CBW-MOVD pairing to handle one variable at a time, use the sign extension capabilities (PSRAW, etc.) within MMX technology for operating on multiple variables. This would result in higher performance as well.
2. Insert another operation that modifies or copies the sign-extended value between the MOVXS/IMUL/CBW instruction and the MOVD instruction as in the example below:
XOR EAX, EAX (or SUB EAX, EAX)
MOVXS AX, BL (or other MOVXS, other IMUL or CBW instruction)
*MOV EAX, EAX
MOVD MM0, EAX

*Note: MOV EAX, EAX is used here as it is fairly generic. Again, EAX can be any 32-bit register.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H24. *Premature Execution of a Load Operation Prior to Exception Handler Invocation*

Problem: This erratum can occur with any of the following situations:

1. If an instruction that performs a memory load causes a code segment limit violation
2. If a waiting floating-point instruction or MMX™ instruction that performs a memory load has a floating-point exception pending
3. If an MMX instruction that performs a memory load and has either CR0.EM =1 (Emulation bit set), or a floating-point Top-of-Stack (FP TOS) not equal to 0, or a DNA exception pending

If any of the above circumstances, occur it is possible that the load portion of the instruction will have executed before the exception handler is entered.

Implication: In normal code execution where the target of the load operation is to write back memory there is no impact from the load being prematurely executed, nor from the restart and subsequent re-execution of that instruction by the exception handler. If the target of the load is to uncached memory that has a system side-effect, restarting the instruction may cause unexpected system behavior due to the repetition of the side-effect.

Workaround: Code which performs loads from memory that has side-effects can effectively workaround this behavior by using simple integer-based load instructions when accessing side-effect memory and by ensuring that all code is written such that a code segment limit violation cannot occur as a part of reading from side-effect memory.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H25. Read Portion of RMW Instruction May Execute Twice

Problem: When the Intel Mobile Celeron processor executes a read-modify-write (RMW) arithmetic instruction, with memory as the destination, it is possible for a page fault to occur during the execution of the store on the memory operand after the read operation has completed but before the write operation completes.

If the memory targeted for the instruction is UC (uncached), memory will observe the occurrence of the initial load before the page fault handler and again if the instruction is restarted.

Implication: This erratum has no effect if the memory targeted for the RMW instruction has no side-effects. If, however, the load targets a memory region that has side-effects, multiple occurrences of the initial load may lead to unpredictable system behavior.

Workaround: Hardware and software developers who write device drivers for custom hardware that may have a side-effect style of design should use simple loads and simple stores to transfer data to and from the device. Then, the memory location will simply be read twice with no additional implications.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H26. Intervening Writeback May Occur During Locked Transaction

Problem: During a transaction which has the LOCK# signal asserted (e.g., a locked transaction), there is a potential for an explicit writeback caused by a previous transaction to complete while the bus is locked. The explicit writeback will only be issued by the processor which has locked the bus, and the lock signal will not be deasserted until the locked transaction completes, but the atomicity of a lock may be compromised by this erratum. Note that the explicit writeback is an expected cycle, and no memory ordering violations will occur. This erratum is, however, a violation of the bus lock protocol.

Implication: A chipset or third-party agent (TPA) which tracks bus transactions in such a way that locked transactions may only consist of a read-write or read-read-write-write locked sequence, with no transactions intervening, may lose synchronization of state due to the intervening explicit writeback. Systems using chipsets or TPAs which can accept the intervening transaction will not be affected.

Workaround: The bus tracking logic of all devices on the system bus should allow for the occurrence of an intervening transaction during a locked transaction.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H27. MC2_STATUS MSR Has Model-Specific Error Code and Machine Check Architecture Error Code Reversed

Problem: The *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, documents that for the MCi_STATUS MSR, bits 15:0 contain the MCA (machine-check architecture) error code field and bits 31:16 contain the model-specific error code field. However, for the MC2_STATUS MSR, these bits have been reversed. For the MC2_STATUS MSR, bits 15:0 contain the model-specific error code field and bits 31:16 contain the MCA error code field.

Implication: A machine check error may be decoded incorrectly if this erratum on the MC2_STATUS MSR is not taken into account.

Workaround: When decoding the MC2_STATUS MSR, reverse the two error fields.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H28. Mixed Cacheability of Lock Variables Is Problematic in MP Systems

Problem: This errata only affects multiprocessor systems where a lock variable address is marked cacheable in one processor and uncacheable in any others. The processors which have it marked uncacheable may stall indefinitely when accessing the lock variable. The stall is only encountered if:

- One processor has the lock variable cached, and is attempting to execute a cache lock.
- The processor which has that address cached has it cached in its L2 only.

Other processors, meanwhile, issue back to back accesses to that same address on the bus.

Implication: MP systems where all processors either use cache locks or consistent locks to uncacheable space will not encounter this problem. If, however, a lock variable's cacheability varies in different processors, and several processors are all attempting to perform the lock simultaneously, an indefinite stall may be experienced by the processors which have it marked uncacheable in locking the variable (if the conditions above are satisfied). Intel has only encountered this problem in focus testing with artificially generated external events. Intel has not currently identified any commercial software which exhibits this problem.

Workaround: Follow a homogenous model for the memory type range registers (MTRRs), ensuring that all processors have the same cacheability attributes for each region of memory; do not use locks whose memory type is cacheable on one processor, and uncacheable on others. Avoid page table aliasing, which may produce a nonhomogenous memory model.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H29. Thermal Sensor May Assert SMBALERT# Incorrectly

Problem: The Intel Celeron Processor Mobile Module has a thermal sensor that monitors the processor core's temperature. Please note that desktop systems could have a similar thermal device. The thermal sensor asserts SMBALERT# if the processor temperature exceeds the temperature limits set in the Alarm Threshold Registers (T_{HIGH} , T_{LOW}). It also sets the corresponding Status Register bits to identify the cause of the interrupt. Figure 1 gives one example of the how the SMBALERT# signal could be used in a system.




Figure 1. An Example of Microcontroller Driven Thermal Management

Implication: There is no system impact from this erratum if temperature polling is used for processor thermal management. If the SMBALERT# interrupt is employed to manage processor thermal sensing, then servicing the false interrupt may result in premature system action depending on the software and hardware implementations used. The rate of the false interrupts is less than the auto-convert rate of the thermal sensor.

Workaround: Three different (mutually exclusive) workarounds are possible:

1. Before servicing an interrupt from the thermal sensor, read and compare the processor thermal reading with the threshold limits (T_{HIGH} or T_{LOW}). Figures 2 and 3 provide basic flowcharts for the implementation of this workaround in an interrupt driven system.
2. If the firmware implemented polls the Status Register only, then before taking any action, re-read the temperature register and do a comparison with the alarm threshold limits (T_{HIGH} or T_{LOW}) to determine if the value is actually still within the temperature window.
3. Use a temperature polling scheme to monitor the processor temperature.

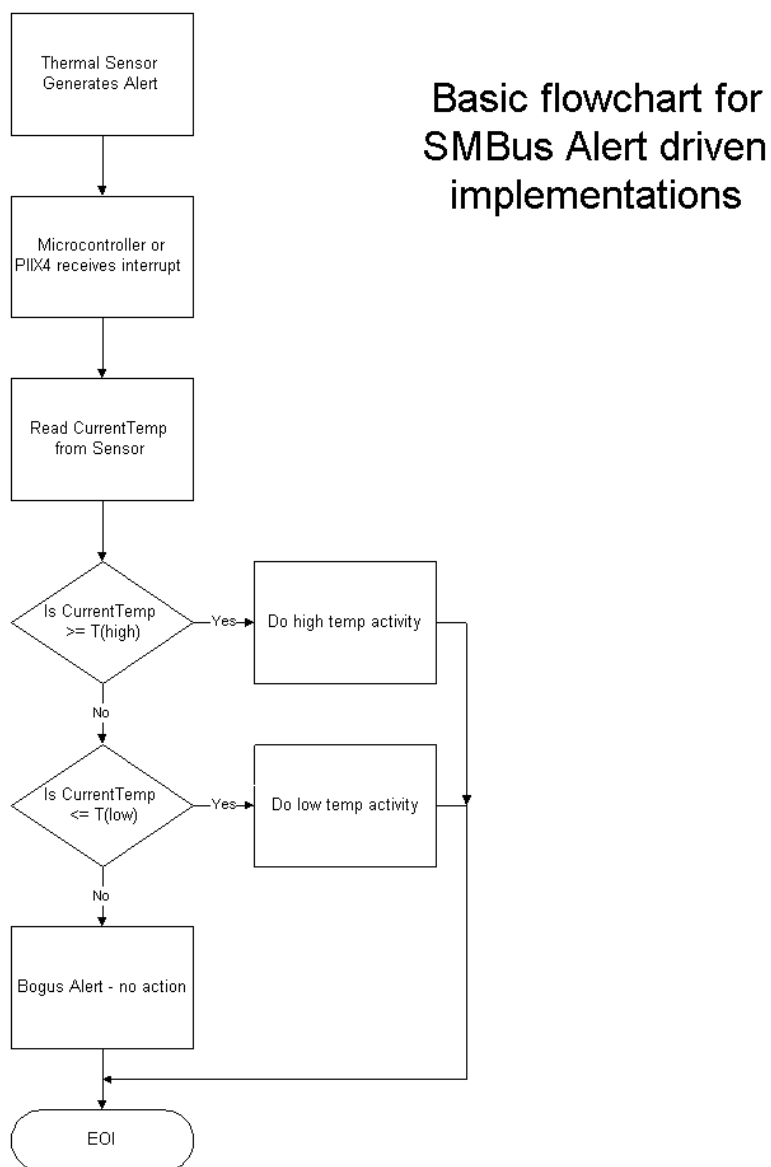


Figure 2. Workaround Flowchart: SMBALERT#-Driven System

Basic flowchart for
system interrupt
driven
implementations

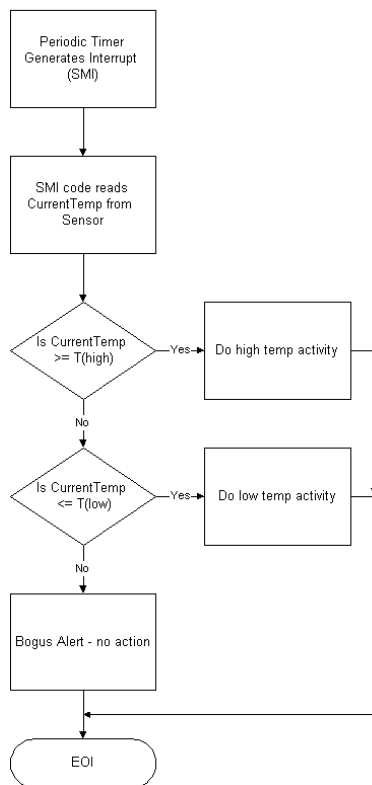


Figure 3. Workaround Flowchart: SMI#-Driven System

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H30. MOV With Debug Register Causes Debug Exception

Problem: When in V86 mode, if a MOV instruction is executed on debug registers, a general-protection exception (#GP) should be generated, as documented in the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, Section 15.2. However, in the case when the general detect enable flag (GD) bit is set, the observed behavior is that a debug exception (#DB) is generated instead.

Implication: With debug-register protection enabled (e.g., the GD bit set), when attempting to execute a MOV on debug registers in V86 mode, a debug exception will be generated instead of the expected general-protection fault.

Workaround: In general, operating systems do not set the GD bit when they are in V86 mode. The GD bit is generally set and used by debuggers. The debug exception handler should check that the exception did not occur in V86 mode before continuing. If the exception did occur in V86 mode, the exception may be directed to the general-protection exception handler.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H31. Upper Four PAT Entries Not Usable With Mode B or Mode C Paging

Problem: The Page Attribute Table (PAT) contains eight entries, which must all be initialized and considered when setting up memory types for the Intel Mobile Celeron processor. However, in Mode B or Mode C paging, the upper four entries do not function correctly for 4-Kbyte pages. Specifically, bit seven of page table entries that translate addresses to 4-Kbyte pages should be used as the upper bit of a three-bit index to determine the PAT entry that specifies the memory type for the page. When Mode B (CR4.PSE = 1) and/or Mode C (CR4.PAE) are enabled, the processor forces this bit to zero when determining the memory type regardless of the value in the page table entry. The upper four entries of the PAT function correctly for 2-Mbyte and 4-Mbyte large pages (specified by bit 12 of the page directory entry for those translations).

Implication: Only the lower four PAT entries are useful for 4KB translations when Mode B or C paging is used. In Mode A paging (4-Kbyte pages only), all eight entries may be used. All eight entries may be used for large pages in Mode B or C paging.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H32. Incorrect Memory Type May be Used When MTRRs Are Disabled

Problem: If the Memory Type Range Registers (MTRRs) are disabled without setting the CR0.CD bit to disable caching, and the Page Attribute Table (PAT) entries are left in their default setting, which includes UC- memory type (PCD = 1, PWT = 0; see the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, for details), data for entries set to UC- will be cached as if the memory type were writeback (WB). Also, if the page tables are set to a memory type other than UC-, then the effective memory type used will be that specified by the page tables and PAT. Any regions of memory normally forced to UC by the MTRRs (such as the VGA video region) may now be incorrectly cached and speculatively accessed.

Even if the CR0.CD bit is correctly set when the MTRRs are disabled and the PAT is left in its default state, then retries and out of order retirement of UC accesses may occur, contrary to the strong ordering expected for these transactions.

Implication: The occurrence of this erratum may result in the use of incorrect data and unpredictable processor behavior when running with the MTRRs disabled. Interaction between the mouse, cursor, and VGA video display leading to video corruption may occur as a symptom of this erratum as well.

Workaround: Ensure that when the MTRRs are disabled, the CR0.CD bit is set to disable caching. This recommendation is described in the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*. If it is necessary to disable the MTRRs, first clear the PAT register before setting the CR0.CD bit, flushing the caches, and disabling the MTRRs to ensure that UC memory type is always returned and strong ordering is maintained.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H33. Misprediction in Program Flow May Cause Unexpected Instruction Execution

Problem: To optimize performance through dynamic execution technology, the P6 architecture has the ability to predict program flow. In the event of a misprediction, the processor will normally clear the incorrect prediction, adjust the EIP to the correct location, and flush out any instructions it may have fetched from the misprediction. In circumstances where a branch misprediction occurs, the correct target of the branch has already been opportunistically fetched into the streaming buffers, and the L2 cycle caused by the evicted cache line is retried by the L2 cache, the processor may fail to flush out the retirement unit before the speculative program flow is committed to a permanent state.

Implication: The results of this erratum may range from no effect to unpredictable application or OS failure. Manifestations of this failure may result in:

- Unexpected values in EIP
- Faults or traps (e.g., page faults) on instructions that do not normally cause faults
- Faults in the middle of instructions
- Unexplained values in registers/memory at the correct EIP

Workaround: It is possible for BIOS code to contain a workaround for this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H34. Data Breakpoint Exception in a Displacement Relative Near Call May Corrupt EIP

Problem: If a data breakpoint is programmed at the memory location where the stack push of a near call is performed, the processor will update the stack and ESP appropriately, but may skip the code at the destination of the call. Hence, program execution will continue with the next instruction immediately following the call, instead of the target of the call.

Implication: The failure mechanism for this erratum is that the call would not be taken; therefore, instructions in the called subroutine would not be executed. As a result, any code relying on the execution of the subroutine will behave unpredictably.

Workaround: Do not program a data breakpoint exception on the stack where the push for the near call is performed.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H35. System Bus ECC Not Functional With 2:1 Ratio

Problem: If a processor is underclocked at a core frequency to system bus frequency ratio of 2:1 and system bus ECC is enabled, the system bus ECC detection and correction will negatively affect internal timing dependencies.

Implication: If system bus ECC is enabled, and the processor is underclocked at a 2:1 ratio, the system may behave unpredictably due to these timing dependencies.

Workaround: All bus agents that support system bus ECC must disable it when a 2:1 ratio is used.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H36. Fault on REP CMPS/SCAS Operation May Cause Incorrect EIP

Problem: If either a General Protection Fault, Alignment Check Fault or Machine Check Exception occur during the first iteration of a REP CMPS or a REP SCAS instruction, an incorrect EIP may be pushed onto the stack of the event handler if all the following conditions are true:

- The event occurs on the initial load performed by the instruction(s)
- The condition of the zero flag before the repeat instruction happens to be opposite of the repeat condition (e.g., REP/REPE/REPZ CMPS/SCAS with ZF = 0 or RENE/REPZ CMPS/SCAS with ZF = 1)
- The faulting micro-op and a particular micro-op of the REP instruction are retired in the retirement unit in a specific sequence

The EIP will point to the instruction following the REP CMPS/SCAS instead of pointing to the faulting instruction.

Implication: The result of the incorrect EIP may range from no effect to unexpected application/OS behavior.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H37. RDMSR or WRMSR To Invalid MSR Address May Not Cause GP Fault

Problem: The RDMSR and WRMSR instructions allow reading or writing of MSRs (Model Specific Registers) based on the index number placed in ECX. The processor should reject access to any reserved or unimplemented MSRs by generating #GP(0). However, there are some invalid MSR addresses for which the processor will not generate #GP(0).

Implication: For RDMSR, undefined values will be read into EDX:EAX. For WRMSR, undefined processor behavior may result.

Workaround: Do not use invalid MSR addresses with RDMSR or WRMSR.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H38. SYSENTER/SYSEXIT Instructions Can Implicitly Load “Null Segment Selector” to SS and CS Registers

Problem: According to the processor specification, attempting to load a null segment selector into the CS and SS segment registers should generate a General Protection Fault (#GP). Although loading a null segment selector to the other segment registers is allowed, the processor will generate an exception when the segment register holding a null selector is used to access memory.

However, the SYSENTER instruction can implicitly load a null value to the SS segment selector. This can occur if the value in SYSENTER_CS_MSR is between FFF8h and FFFBh when the SYSENTER instruction is executed. This behavior is part of the SYSENTER/SYSEXIT instruction definition; the content of the SYSTEM_CS_MSR is always incremented by 8 before it is loaded into the SS. This operation will set the null bit in the segment selector if a null result is generated, but it does not generate a #GP on the SYSENTER instruction itself. An exception will be generated as expected when the SS register is used to access memory, however.

The SYSEXIT instruction will also exhibit this behavior for both CS and SS when executed with the value in SYSENTER_CS_MSR between FFF0h and FFF3h, or between FFE8h and FFEb, inclusive.

Implication: These instructions are intended for operating system use. If this erratum occurs (and the OS does not ensure that the processor never has a null segment selector in the SS or CS segment registers), the processor's behavior may become unpredictable, possibly resulting in system failure.

Workaround: Do not initialize the SYSTEM_CS_MSR with the values between FFF8h and FFFBh, FFF0h and FFF3h, or FFE8h and FFEb before executing SYSENTER or SYSEXIT.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H39. PRELOAD Followed by EXTEST Does Not Load Boundary Scan Data

Problem: According to the IEEE 1149.1 Standard, the EXTEST instruction would use data “typically loaded onto the latched parallel outputs of boundary-scan shift-register stages using the SAMPLE/PRELOAD instruction prior to the selection of the EXTEST instruction.” As a result of this erratum, this method cannot be used to load the data onto the outputs.

Implication: Using the PRELOAD instruction prior to the EXTEST instruction will not produce expected data after the completion of EXTEST.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H40. Far Jump to New TSS With D-bit Cleared May Cause System Hang

Problem: A task switch may be performed by executing a far jump through a task gate or to a new Task State Segment (TSS) directly. Normally, when such a jump to a new TSS occurs, the D-bit (which indicates that the page referenced by a Page Table Entry (PTE) has been modified) for the PTE which maps the location of the previous TSS will already be set and the processor will operate as expected. However, if the D-bit is clear at the time of the jump to the new TSS, the processor will hang.

Implication: If an OS is used which can clear the D-bit for system pages, and which jumps to a new TSS on a task switch, then a condition may occur which results in a system hang. Intel has not identified any commercial software which may encounter this condition; this erratum was discovered in a focused testing environment.

Workaround: Ensure that OS code does not clear the D-bit for system pages (including any pages that contain a task gate or TSS). Use task gates rather than jumping to a new TSS when performing a task switch.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H41. Incorrect Chunk Ordering May Prevent Execution of the Machine Check Exception Handler After BINIT#

Problem: If a catastrophic bus error is detected which results in a BINIT# assertion, and the BINIT# assertion is propagated to the processor's L2 cache at the same time that data is being sent to the processor, then the data may become corrupted in the processor's L1 cache.

Implication: Since BINIT# assertion is due to a catastrophic event on the bus, the corrupted data will not be used. However, it may prevent the processor from executing the Machine Check Exception (MCE) handler, causing the system to hang.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.



H42. Resume Flag May Not Be Cleared After Debug Exception

Problem: The Resume Flag (RF) is normally cleared by the processor after executing an instruction which causes a debug exception (#DB). In the process of determining whether the RF needs to be cleared after executing the instruction, the processor uses an internal register containing stale data. The stale data may unpredictably prevent the processor from clearing the RF.

Implication: If this erratum occurs, further debug exceptions will be disabled.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H43. Processor May Return Invalid Parameters on Execution of the CPUID Instruction

Problem: The Mobile Module based Intel Celeron processor with on-die L2 cache may return an incorrect "Maximum CPUID Input Value." The Mobile Module based Intel Celeron processor with on-die L2 cache is specified to return a value of 2 in the EAX register when the CPUID instruction is executed with EAX=0; however, this erratum may result in the value of 3 being returned in EAX. It is also possible that bit 18 of the EDX register will be set to a 1 when CPUID is executed with EAX=1. This bit is defined as reserved for the Mobile Module based Intel Celeron processor with on-die L2 cache, but is expected to be set to zero. If CPUID were to be executed on the Mobile Module based Intel Celeron processor with on-die L2 cache with EAX=3 the processor should return the cache parameters in the integer registers (EAX, EBX, ECX, EDX); however, the processors affected by this erratum will return undefined values in the integer registers.

Implication: Intel has not seen any software failures as a result of this erratum; however, since software written for the Mobile Module based Intel Celeron processor with on-die L2 cache will not be expecting to see a Maximum CPUID Input Value greater than 2, it is not possible to predict how software will behave on processors with this erratum. Software using the fact that bit 18 in the feature flags is set, to determine the presence of the Pentium® III processor serial number feature, without also verifying that it is executing on a Pentium III processor, may incorrectly believe that the Mobile Module based Intel Celeron processor with on-die L2 cache, support the processor serial number feature. However, the values the CPUID instruction returns when CPUID is executed with EAX=3 will not be the processor serial number and will be undefined.

Workaround: It is possible for BIOS code to contain a workaround for this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H44. Internal Cache Protocol Violation May Cause System Hang

Problem: An Intel Mobile Celeron processor based system may hang due to an internal cache protocol violation. During multiple transactions targeted at the same cacheline, there exists a small window of time such that the processor's internal timings align to create a livelock situation. The scenario, which results in the erratum, is summarized below:

Scenario:

1. A snoopable transaction is issued to address A. This snoopable transaction can be issued by the processor or the chipset.
2. The snoopable transaction hits a modified line in the processor's L1 data cache.
3. The processor issues two code fetches from the L2 cache before the snoopable transaction reaches the top of the In-Order Queue and before the snoopable transaction's modified L1 cache line containing address A is brought out on the system bus.
4. At the same time, a locked access to the L1 cache occurs.

Implication: An Intel Mobile Celeron processor may cause a system to hang if the above listed sequence of events occur. The probability of encountering this erratum increases with I/O queue depth greater than four.

Workaround: It is possible for the BIOS code to contain a workaround for this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H45. GP# Fault on WSMSR to ROB_CR_BKUPTMPDR6

Problem: Writing a '1' to unimplemented bit(s) in the ROB_CR_BKUPTMPDR6 MSR (offset 1E0H) will result in a general protection fault (GP#).

Implication: The normal process used to write an MSR is to read the MSR using RDMSR, modify the bit(s) of interest, and then to write the MSR using WRMSR. Because of this erratum, this process may result in a GP# fault when used to modify the ROB_CR_BKUPTMPDR6 MSR.

Workaround: When writing to ROB_CR_BKUPTMPDR6 all unimplemented bits must be '0.' Implemented bits may be set as '0' or '1' as desired.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.



H46. Machine Check Exception May Occur Due to Improper Line Eviction in the IFU

Problem: The Intel Mobile Celeron processor is designed to signal an unrecoverable Machine Check Exception (MCE) as a consistency checking mechanism. Under a complex set of circumstances involving multiple speculative branches and memory accesses there exists a one cycle long window in which the processor may signal a MCE in the Instruction Fetch Unit (IFU) because instructions previously decoded have been evicted from the IFU. The one cycle long window is opened when an opportunistic fetch receives a partial hit on a previously executed but not as yet completed store resident in the store buffer. The resulting partial hit erroneously causes the eviction of a line from the IFU at a time when the processor is expecting the line to still be present. If the MCE for this particular IFU event is disabled, execution will continue normally.

Implication: While this erratum may occur on a system with any number of processors, the probability of occurrence increases with the number of processors. If this erratum does occur, a machine check exception will result. Note systems that implement an operating system that does not enable the Machine Check Architecture will be completely unaffected by this erratum, e.g., Windows* 95 and Windows 98.

Workaround: It is possible for BIOS code to contain a workaround for this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H47 Lower Bits of SMRAM SMBASE Register Cannot Be Written With an ITP

Problem: The System Management Base (SMBASE) register (7EF8H) stores the starting address of the System Management RAM (SMRAM). This register is used by the processor when it is in System Management Mode (SMM), and its contents serve as the memory base for code execution and data storage. The 32-bit SMBASE register can normally be programmed to any value. When programmed with an In-Target Probe (ITP), however, any attempt to set the lower 11 bits of SMBASE to anything other than zeros via the WRMSR instruction will cause the attempted write to fail.

Implication: When set via the ITP, any attempt to relocate SMRAM space must be made with 2 Kbyte alignment.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

Task Switch Caused by Page Fault May Cause Wrong PTE H48. and PDE Access Bit to be Set

Problem: If an operating system services a page fault through a Task Gate, the access bit may be set for an incorrect page table/directory entry.

Implication: An operating system that uses a Task Gate for its page fault handler may encounter this erratum. The effect of the erratum depends on the alignment of the Task State Segment (TSS), and ranges from no anomalous behavior to unexpected errors. Intel is not aware of any commercial operating systems which use a Task Gate to handle page faults. Task gates used for other purposes (NMI, Machine Check, or Double Fault) do not cause this erratum.

Workaround: The operating system may alternately use an Interrupt Gate or a Call Gate rather than a Task Gate.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

Cross Modifying Code Operations on a Jump Instruction H49. May Cause a General Protection Fault

Problem: The act of one processor writing data into the currently executing code segment of a second processor with the intent of having the second processor execute that data as code is called Cross-Modifying Code (XMC). Software using XMC to modify the offset of an execution transfer instruction (i.e., Jump, Call etc.) without a synchronizing instruction may cause a General Protection Fault (GPF) when the offset splits a cache line boundary.

Implication: Any application creating a (GPF) would be terminated by the operating system.

Workaround: Programmers should use the cross modifying code synchronization algorithm as detailed in Volume 3 of the Intel Architecture Software Developer's Manual, section 7.1.3, in order to avoid this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.



H50. Deadlock May Occur Due To Illegal-Instruction/Page-Miss Combination

Problem: Intel's 32-bit Instruction Set Architecture (ISA) utilizes most of the available op-code space; however some byte combinations remain undefined and are considered illegal instructions. Intel processors detect the attempted execution of illegal instructions and signal an exception. This exception is handled by the operating system and/or application software.

Under a complex set of internal and external conditions involving illegal instructions, a deadlock may occur within the processor. The necessary conditions for the deadlock involve:

1. The illegal instruction is executed.
2. Two page table walks occur within a narrow timing window coincident with the illegal instruction.

Implication: The illegal instructions involved in this erratum are unusual and invalid byte combinations that are not useful to application software or operating systems. These combinations are not normally generated in the course of software programming, nor are such sequences known by Intel to be generated in commercially available software and tools. Development tools (compilers, assemblers) do not generate this type of code sequence, and will normally flag such a sequence as an error. If this erratum occurs, the processor deadlock condition will occur and result in a system hang. Code execution cannot continue without a system RESET.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

H51. FLUSH# Assertion Following STPCLK# May Prevent CPU Clocks From Stopping

Problem: If FLUSH# is asserted after STPCLK# is asserted, the cache flush operation will not occur until after STPCLK# is de-asserted. Furthermore, the pending flush will prevent the processor from entering the Sleep state, since the flush operation must complete prior to the processor entering the Sleep state.

Implication: Following SLP# assertion, processor power dissipation may be higher than expected. Furthermore, if the source to the processor's input bus clock (BCLK) is removed, normally resulting in a transition to the Deep Sleep state, the processor may shutdown improperly. The ensuing attempt to wake up the processor will result in unpredictable behavior and may cause the system to hang.

Workaround: For systems that use the FLUSH# input signal and Deep Sleep state of the processor, ensure that FLUSH# is not asserted while STPCLK# is asserted.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

DOCUMENTATION CHANGES

The Documentation Changes listed in this section applies to the following documents:

- *Intel® Mobile Celeron™ Processor in Micro-PGA and BGA packages at 466 MHz, 433 MHz, 400 MHz, 366 MHz, 333 MHz, 300 MHz, and 266 MHz* datasheet
- *Intel® Mobile Celeron™ Processor: Mobile Module MMC-1*
- *Intel® Mobile Celeron™ Processor: Mobile Module MMC-2*
- *Intel Architecture Software Developer's Manual, Volumes 1, 2 and 3*

All Documentation Changes will be incorporated into a future version of the appropriate Intel Celeron processor documentation.

H1. SMBus Data Setup Time

The SMBus data signal has a min setup time (min Tsu: DAT) of 250 ns, as published in the *System Management Bus Specification*, Rev 1.0. This parameter is currently specified to 800 ns for the Mobile Pentium II processor (mini-cartridge) and Intel Pentium II Processor Mobile Module products (see table below).

SMBus Data Min Setup

Parameter	Description	Spec	Change to
min Tsu: DAT	min SMBus Data Setup Time	250 ns	800 ns

H2. Intel® Celeron™ Processor Mobile Module Connector 1 (MMC-1) Northbridge Correction

In the *Intel® Mobile Celeron™ Processor: Mobile Module MMC-1* datasheet, the northbridge component was incorrectly identified as the 443BX. The correct northbridge component for MMC-1 is the 443DX.

H3. VR_ON and VRPWRGD Specification Correction for MMC-1 and MMC-2

In the *Intel® Mobile Celeron™ Processor: Mobile Module MMC-1* datasheet and the *Intel® Mobile Celeron™ Processor: Mobile Module MMC-2* datasheet, the VR_ON and VR_PWRGD signal specifications are incorrectly identified as V_3S in all of the above datasheets. The signal specifications should be pulled up to V_3.

H4. STR and POS Measurement Correction for MMC-1 and MMC-2

In the *Intel® Mobile Celeron™ Processor: Mobile Module MMC-1* datasheet and the *Intel® Mobile Celeron™ Processor: Mobile Module MMC-2* datasheet, the POS and STR values are incorrectly shown as measured in

Amps in all of the above datasheets. The POS and STR values should be measured in Watts. The actual number values are correct.

H5. L2 Cache Correction for MMC-1 and MMC-2

In the *Intel® Mobile Celeron™ Processor: Mobile Module MMC-1* datasheet and the *Intel® Mobile Celeron™ Processor: Mobile Module MMC-2* datasheet, V_3S is incorrectly identified as the power source for the L2 Cache circuit. V_3S is not the power source for the L2 Cache circuit.

H6. Die Height (A₂) in Surface-Mount BGA1 Package Specification

In the *Intel® Mobile Celeron™ Processor in BGA Package at 300 MHz and 266 MHz* datasheet, the original documentation of Min at 1.23 mm and Max at 1.38 mm is incorrect. The correct documentation is 0.841 mm REF.

Table 0.1 Surface-Mount BGA1 Package Specifications

Symbol	Parameter	Min	Max	Unit
A	Overall Height, as delivered	2.29	2.79	mm
A ₁	Ball Height, as delivered	0.76	1.10	mm
A ₂	Die Height	0.841 REF		mm
b	Ball Diameter	0.78 REF		mm
D	Package Width	30.850	31.150	mm
D ₁	Die Width	10.36 REF		mm
E	Package Length	34.850	35.150	mm
e	Ball Pitch	1.270		mm
E ₁	Die Length	17.36 REF		mm
K	Keepout Outline to Edge of Substrate	5 REF		mm
K ₁	Keepout Outline to Edge of Substrate at Corner	7 REF		mm
N	Ball Count	615		each
S ₁	Outer Ball Center to Short Edge of Substrate	1.625 REF		mm
S ₂	Outer Ball Center to Long Edge of Substrate	0.895 REF		mm
P _{DIE}	Allowable Pressure on the Die for Thermal Solution	—	689	kPa
W	Package Weight	3.71	4.52	grams

H7. MTRR Initialization Clarification

The following sentence should be added to the end of the first paragraph of Section 9.12.5 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*: "The MTRRs must be disabled prior to initialization or modification."

H8. STPCLK# Pin Definition

The *Intel® Mobile Celeron™ Processor in BGA Package at 333 MHz, 300 MHz and 266 MHz* datasheet (Order Number 245112-002) has an incorrect definition of the STPCLK# pin. This document incorrectly states,

The processor continues to snoop bus transactions and *service interrupts* while in Stop Grant state. When STPCLK# is deasserted, the processor restarts its internal clock to all units and resumes execution.

It should state,

The processor continues to snoop bus transactions and *may latch interrupts* while in Stop Grant state. When STPCLK# is deasserted, the processor restarts its internal clock to all units, resumes execution, *and services any pending interrupts*.

H9. RESET# Pin Definition

The *Intel® Mobile Celeron™ Processor in BGA Package at 333 MHz, 300 MHz and 266 MHz* datasheet (Order Number 245112-002) has an incorrect definition of the RESET# pin. This document incorrectly states:

RESET# must remain active for one microsecond for a "warm" Reset; for a Power-on Reset, RESET# must stay active for at least one millisecond after VCCCORE and CLK have reached their proper specifications.

It should state,

For a Power-on Reset, RESET# must stay active for at least one millisecond after VCCCORE and CLK have reached their proper specifications.

H10. BGA1 Package Mechanical Specifications

The following table and figures show the updated mechanical specifications for the BGA1 package, listed as Table 5.2, Figure 5.3 and 5.4, respectively, in the *Intel® Mobile Celeron™ Processor in BGA Package at 333 MHz, 300 MHz and 266 MHz* datasheet.

Table 5.2 Surface-Mount BGA1 Package Specifications

Symbol	Parameter	Min	Max	Unit
A	Overall Height, as delivered	2.29	2.79	mm
A ₁	Substrate Height, as delivered	1.50 REF		mm
A ₂	Die Height	0.854 REF ¹		mm
B	Ball Diameter	0.78 REF		mm
D	Package Width	30.85	31.15	mm
D ₁	Die Width	10.36 REF		mm
E	Package Length	34.85	35.15	mm
e	Ball Pitch	1.27		mm
E ₁	Die Length	17.36 REF		mm
N	Ball Count	615		each
S ₁	Outer Ball Center to Short Edge of Substrate	1.625 REF		mm
S ₂	Outer Ball Center to Long Edge of Substrate	0.895 REF		mm
P _{DIE}	Allowable Pressure on the Die for Thermal Solution	—	689	kPa
W	Package Weight	3.71	4.52	grams

NOTES:

1. The number shown here is different from that in the 244444-004 version of this specification update. The maximum reference number is used.

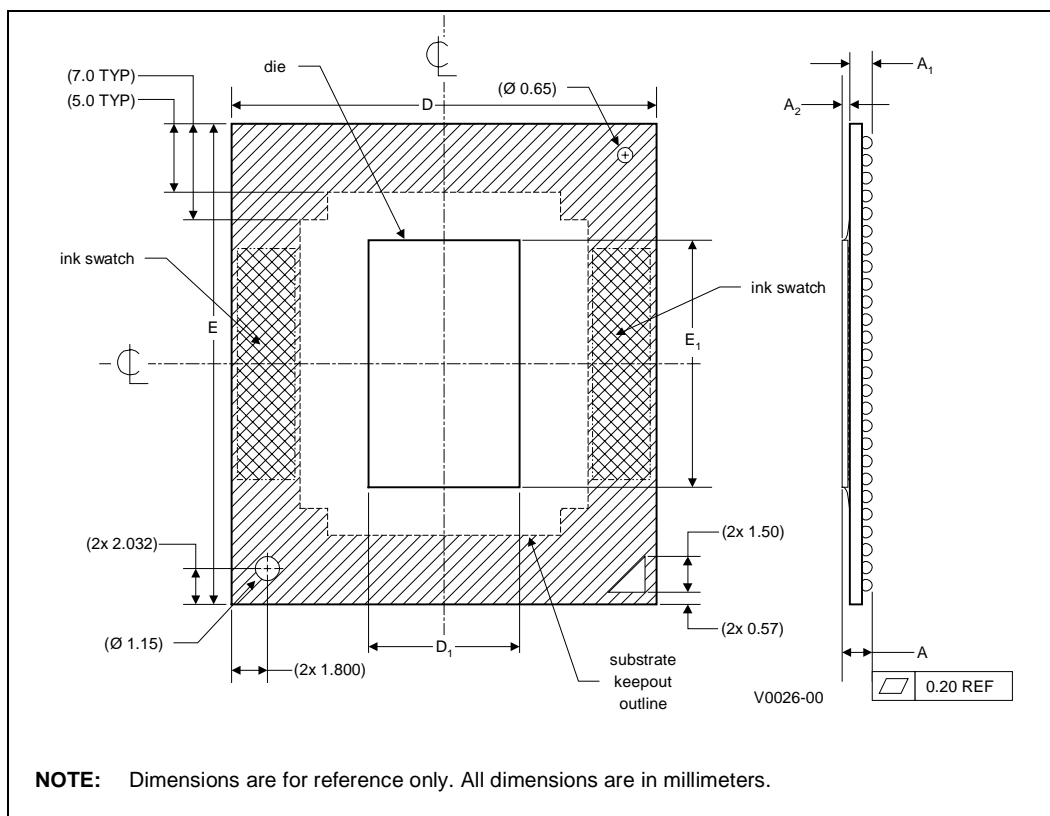



Figure 5.3. Surface-Mount BGA1 Package—Top and Side View



V0025-01

Figure 5.4. Surface-Mount BGA1 Package—Bottom View

H11. Generic Clock Waveform Figure

Figure 3.3 in the *Intel® Mobile Celeron™ Processor in μ PGA1 and BGA1 Packages at 366 MHz, 333 MHz, 300 MHz, and 266 MHz* datasheet (Order Number 245112-003) has been updated. The following figure is to replace the one in the above document.

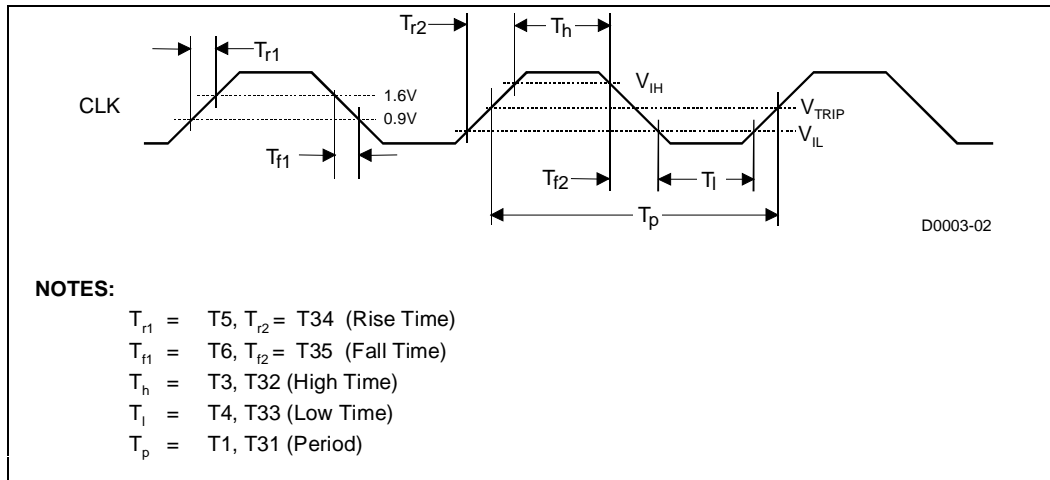


Figure 3.3 Generic Clock Waveform

H12. Invalidating Caches and TLBs

Section 2.6.4 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide* incorrectly states:

The INVD (invalidate cache with no writeback) instruction invalidates all data and instruction entries in the internal caches and TLBs and sends a signal to the external caches indicating that they should be invalidated also.

It should state:

The INVD (invalidate cache with no writeback) instruction invalidates all data and instruction entries in the internal caches and sends a signal to the external caches indicating that they should be invalidated also.

H13. Handling of Self-Modifying and Cross-Modifying Code

Section 7.1.3 paragraph 1. of the *Intel Architecture Software Developer's Manual Vol 3: System Programming* incorrectly states:

The act of a processor writing data into a currently executing code segment with the intent of executing that data as code is called **self-modifying code**. Intel Architecture processors exhibit model-specific behavior when executing self-modified code, depending upon how far ahead of the current execution pointer the code has been modified. As processor architectures become more complex and start to speculatively execute code ahead of the retirement point (as in the P6 family processors), the rules regarding which code should execute, pre- or post-modification, become blurred. To write self-modifying code and ensure that it is compliant with current and future Intel Architectures one of the following two coding options **should** be chosen.

It should state:

The act of a processor writing data into a currently executing code segment with the intent of executing that data as code is called **self-modifying code**. Intel Architecture processors exhibit model-specific behavior when executing self-modified code, depending upon how far ahead of the current execution pointer the code has been modified. As processor architectures become more complex and start to speculatively execute code ahead of the retirement point (as in the P6 family processors), the rules regarding which code should execute, pre- or post-modification, become blurred. To write self-modifying code and ensure that it is compliant with current and future Intel Architectures one of the following two coding options **must** be chosen.

Section 7.1.3 paragraph 6. of the *Intel Architecture Software Developer's Manual Vol 3: System Programming* incorrectly states:

The act of one processor writing data into the currently executing code segment of a second processor with the intent of having the second processor execute that data as code is called **cross-modifying code**. As with self-modifying code, Intel Architecture processors exhibit model-specific behavior when executing cross-modifying code, depending upon how far ahead of the executing processors current execution pointer the code has been modified. To write cross-modifying code and insure that it is compliant with current and future Intel Architectures, the following processor synchronization algorithm **should** be implemented.

It should state:

The act of one processor writing data into the currently executing code segment of a second processor with the intent of having the second processor execute that data as code is called **cross-modifying code**. As with self-modifying code, Intel Architecture processors exhibit model-specific behavior when executing cross-modifying code, depending upon how far ahead of the executing processors current execution pointer the code has been modified. To write cross-modifying code and insure that it is compliant with current and future Intel Architectures, the following processor synchronization algorithm **must** be implemented.

H14. Machine Check Architecture Initialization of MCI_STATUS Registers

Section 12.5, the last paragraph of the *Intel Architecture Software Developer's Manual Vol. 3: System Programming Guide* incorrectly states:

The processor can write valid information (such as an ECC error) into the MCI_STATUS registers while it is being powered up. As part of the initialization of the MCE exception handler, software might examine all the MCI_STATUS registers and log the contents of them, then rewrite them all to zeroes. This procedure is not included in the initialization pseudocode in Example 12-1.

It should state:

The processor can write valid information (such as an ECC error) into the MCI_STATUS registers while it is being powered up. As part of the initialization of the MCE exception handler, software might examine all the MCI_STATUS registers and log the contents of them, then rewrite them all to zeroes. Following power cycling, the MCI_STATUS registers are not guaranteed to have valid data until after the registers are initially cleared to all zeroes by software. This procedure is not included in the initialization pseudocode in Example 12-1.

H15. LOCK# Signal Prefix Operands

Page 3-273, the first sentence of the third paragraph of the *Intel Architecture Software Developer's Manual, Volume 2: Instruction Set Reference* states:

The LOCK prefix can be prepended only to the following instructions and to those forms of the instructions that use a memory operand: ADD, ADC, AND, ...

It should state:

The LOCK prefix can be prepended only to the following instructions and to those forms of the instructions that use a memory operand as the destination operand only: ADD, ADC, AND, ...

If the LOCK prefix is used with the memory operand as the source operand then an Invalid Opcode (Undefined Opcode), #UD, can occur.

SPECIFICATION CLARIFICATIONS

The Specification Clarifications listed in this section apply to the following documents:

- *Intel® Mobile Celeron™ Processor in Micro-PGA and BGA Packages at 466 MHz, 433 MHz, 400 MHz, 366 MHz, 333 MHz, 300 MHz, and 266 MHz datasheet*
- *Intel® Mobile Celeron™ Processor: Mobile Module MMC-1*
- *Intel® Mobile Celeron™ Processor: Mobile Module MMC-2*

All Specification Clarifications will be incorporated into a future version of the appropriate Intel Celeron processor documentation.

H1. Shipping Container Maximum Temperature Rating for BGA1

In the *Intel® Mobile Celeron™ Processor in BGA Package at 300 MHz and 266 MHz* datasheet, note 1 is added to Table 3.5 of the datasheets to include the maximum shipping container temperature rating of 65° C. The following is the updated Table 3.5.

Table 3.5. Mobile Celeron™ Processor Absolute Maximum Ratings

Symbol	Parameter	Min	Max	Unit	Notes
T _{Storage}	Storage Temperature	−40	85	°C	1
V _{CC(Abs)}	Supply Voltage with respect to V _{SS}	−0.5	3.0	V	
V _{CCP}	CMOS Reference Voltage with respect to V _{SS}	−0.3	3.0	V	
V _{IN}	GTL+ Buffer DC Input Voltage with respect to V _{SS}	−0.3	V _{CC} + 0.7	V	2
V _{IN25}	2.5 V Buffer DC Input Voltage with respect to V _{SS}	−0.3	3.3	V	3

NOTES:

1. The shipping container is only rated for 65° C.
2. Parameter applies to the Low Power GTL+ signal groups only.
3. Parameter applies to CMOS, Open-Drain, APIC and TAP bus signal groups only.

H2. ESD for MMC-1 and MMC-2

In the *Intel® Mobile Celeron™ Processor: Mobile Module MMC-1* datasheet and the *Intel® Mobile Celeron™ Processor: Mobile Module MMC-2* datasheet, ESD is a nonpowered test of the module for noncatastrophic failure only. The module is tested at 2 KV and then is inserted in a system for a functional test.

H3. Bulk Capacitance requirements for the system electronics

This is a clarification to the *Intel® Celeron Processor: Mobile Module MMC-1 (Order Number #245101)*, the *Intel® Mobile Celeron Processor: Mobile Module MMC-2 (Order Number #245102)*:

Placement of sufficient bulk capacitance on the system electronics board is critical to the operation of the Intel Celeron processor mobile module and to ensure the system design can accommodate future high frequency modules. Intel has provided the maximum possible bulk capacitance on the module. However, in order to achieve proper filtering and in-rush current protection, it is imperative that additional filtering be provided on the system electronics board. Table 1 details the bulk capacitance requirements for the system electronics. Also, take special note of the voltage rating requirement for the capacitors on each respective voltage rail.

Note that the rating requirement and note 5 has been changed from previous versions of this table in the *Celeron Processor Mobile Module datasheet*. Future revisions of this specification will include these new changes.

Table 1. Bulk Capacitance

Power Plane	Capacitance Requirements	Max Total ESR	Ripple Current	Rating
V_DC	100 uf ⁵ , 0.1 uf, 0.01 uf ¹	20 mΩ	4.5A-7.5A ³ 3A ~ 5A ⁴	20% tolerance at 35V
V_5	100 uf ⁵ , 0.1 uf, 0.01 uf ¹	100 mΩ	1A	20% tolerance at 10V
V_3	470 uf ⁵ , 0.1 uf, 0.01 uf ¹	100 mΩ	1A	20% tolerance at 6V
V_3S	100 uf ⁵ , 0.1 uf, 0.01 uf ¹	100 mΩ	N/A	20% tolerance at 6V
VCC_AGP	22 uf ⁵ , 0.1 uf, 0.01 uf ¹	100 mΩ	1A	20% tolerance at 6V
V_CPUPU	2.2 uf, 8200 pf ¹	N/A	N/A	20% tolerance at 6V
V_CLK	10 uf ⁵ , 8200 pf ²	N/A	N/A	20% tolerance at 6V

NOTES:

1. Placement of above capacitance requirements should be located near the connector.
2. V_CLK filtering should be located next to the system clock synthesizer.
3. Ripple current specification depends on V_DC input for the Geyserville module.
4. Ripple current specification depends on V_DC input for the single frequency module.
5. If Tantalum Capacitors are used, a 50% voltage de-rating practice must be observed (for example, a 5.0-V rail requires a 10.0-V rated capacitor).

SPECIFICATION CHANGES

The Specification Changes listed in this section apply to the following documents:

- *Intel® Mobile Celeron™ Processor in Micro-PGA and BGA Packages at 466 MHz, 433 MHz, 400 MHz, 366 MHz, 333 MHz, 300 MHz, and 266 MHz* datasheet
- *Intel® Mobile Celeron™ Processor: Mobile Module MMC-1*
- *Intel® Mobile Celeron™ Processor: Mobile Module MMC-2*

All Specification Changes will be incorporated into a future version of the appropriate Intel Mobile Celeron processor documentation.

H1. *ICC,SG, ICC,QS, and ICC,DSLP Maximum Specifications for BGA1 and Mini-Cartridge*

In the *Intel® Mobile Celeron™ Processor in BGA Package at 300 MHz and 266 MHz* datasheet, maximum $I_{CC,SG}$, $I_{CC,QS}$, and $I_{CC,DSLP}$ specifications have been increased by 250 mA, to 1190 mA, 880 mA and 650 mA respectively. The following is the updated Table 3.6 of the datasheets.

Table 3.6. Mobile Celeron™ Processor Power Specifications¹

$T_{CASE} = 0$ to $T_{CASE,max}$; $V_{CC} = 1.6\text{ V} \pm 135\text{ mV}$; $V_{CCP} = 1.8\text{ V} \pm 90\text{ mV}$						
Symbol	Parameter	Min	Typ	Max	Unit	Notes
V_{CC}	V_{CC} of core logic for regular voltage processors	1.465	1.6	1.735	V	$\pm 135\text{ mV}$
$V_{CC,LP}$	V_{CC} when $I_{CC} < 300\text{ mA}$	1.465	1.6	1.805	V	+205/-135 mV ²
V_{CCP}	V_{CC} for CMOS voltage references	1.71	1.8	1.89	V	1.8 V $\pm 90\text{ mV}$
I_{CC}	I_{CC} for V_{CC} at core @ 300 MHz frequency @ 266 MHz			7.49 6.63	A A	5
I_{CCP}	Current for V_{CCP}			75	mA	3, 4, 5
$I_{CC,SG}$	Processor Stop Grant and Auto Halt current			1190	mA	5
$I_{CC,QS}$	Processor Quick Start and Sleep current			880	mA	5
$I_{CC,DSLP}$	Processor Deep Sleep leakage current			650	mA	5
dI_{CC}/dt	V_{CC} power supply current slew rate			20	A/ μs	6, 7

NOTES:

1. Unless otherwise noted, all specifications in this table apply to all processor frequencies.
2. A higher $V_{CC,MAX}$ is allowed when the processor is in a low power state to enable high efficiency, low current modes in the power regulator.
3. I_{CCP} is the current supply for the CMOS voltage references.

4. Not 100% tested. Specified by design/characterization.
5. $I_{CCx,max}$ specifications are specified at $V_{CC,max}$, $V_{CCP,max}$ and 100° C and under maximum signal loading conditions. $I_{CCx,max}$ specifications are not specified at $V_{CC,LP,max}$, if that voltage specification is used then slightly higher currents can be expected.
6. Based on simulations and averaged over the duration of any change in current. Use to compute the maximum inductance and reaction time of the voltage regulator. This parameter is not tested.
7. Maximum values specified by design/characterization at nominal V_{CC} and V_{CCP} .

H2. Temperature note for thermal power specifications

In the *Intel® Mobile Celeron™ Processor in BGA Package at 333 MHz, 300 MHz and 266 MHz* datasheet, Table 6.1 TDP notes should state “at 100° C^{2, 3}” instead of “at 50° C^{2, 3}”. Table Note 3 should state, “3. Not 100% tested or guaranteed. The power specifications are composed of the current of the processor on the various voltage planes. These currents are measured and specified at high temperature in Section 3.5. The 50° C power specifications are determined by characterization of the processor currents at higher temperatures.”